arXiv:1403.5715v2 [cs.CR] 27 Mar 2014

Mining Attribute-Based Access Control Policies from Logs*

Zhongyuan Xu and Scott D. Stoller
Department of Computer Science, Stony Brook University, USA

March 28, 2014

Abstract

Attribute-based access control (ABAC) provides a high level of flexibility that promotes security and information sharing. ABAC policy mining algorithms have potential to significantly reduce the cost of migration to ABAC, by partially automating the development of an ABAC policy from information about the existing access-control policy and attribute data. This paper presents an algorithm for mining ABAC policies from operation logs and attribute data. To the best of our knowledge, it is the first algorithm for this problem.

1 Introduction

ABAC is becoming increasingly important as security policies become more dynamic and more complex. In industry, more and more products support ABAC, using a standardized ABAC language such as XACML [XAC] or a vendor-specific ABAC language. In government, the Federal Chief Information Officer Council called out ABAC as a recommended access control model [Fed11, HFK+13]. ABAC allows "an unprecedented amount of flexibility and security while promoting information sharing between diverse and often disparate organizations" [HFK+13]. ABAC overcomes some of the problems associated with RBAC, notably role explosion [HFK+13], which makes RBAC policies large and hard to manage.

ABAC promises long-term cost savings through reduced management effort, but manual development of an initial policy can be difficult [BM13] and expensive [HFK+13]. *Policy mining* algorithms promise to drastically reduce the cost of migrating to ABAC, by partially automating the process.

Role mining, i.e., mining of RBAC policies, is an active research area (e.g., [KSS03, SS05, VAW06, VAG07, ZRE07, CDPO08, FBB08, GVA08, LVA08, VAGA08, CDPOV09, MLL+09, FBB10, MLL10, MCL+10, MLQ+10, TJ10, VAWG10, CDV12, LVAH12, MPC12, VVAC12, XS12, FBB13, MSAV13, XS13b, ZCG+13]) and a currently relatively small (about $70 million) but rapidly growing commercial market segment [HCBC12]. Some form of role mining is supported by security products from companies of all sizes, from IT giants (e.g., IBM Tivoli Access Manager, Oracle Role Manager, CA Technologies GovernanceMinder) to small start-ups (e.g., Bay31 Role Designer). In contrast, there is, so far, relatively little work on ABAC policy mining. We recently developed an algorithm to mine an ABAC policy from an ACL policy or RBAC policy [XS13a]; to the best of our knowledge, it is the first algorithm for that problem.

However, an ACL policy or RBAC policy might not be available, e.g., if the current access control policy is encoded in a program or is not enforced by a computerized access control mechanism. An alternative source of information about the current access control policy is operation logs, or "logs" for short. Many software systems produce logs, e.g., for auditing, accounting, and accountability purposes. Molloy, Park, and

*This material is based upon work supported in part by NSF under Grant CNS-0831298.

Chari proposed the idea of mining policies from logs and developed algorithms for mining RBAC policies from logs [MPC12].

The main challenge is that logs generally provide incomplete information about entitlements (i.e., granted permissions). Specifically, logs provide only a lower bound on the entitlements. Therefore, the generated policy should be allowed to include *over-assignments*, i.e., entitlements not reflected in the logs.

This paper presents an algorithm for mining ABAC policies from logs and attribute data. To the best of our knowledge, it is the first algorithm for this problem. It is based on our algorithm for mining ABAC policies from ACLs [XS13a]. At a high level, the algorithm works as follows. It iterates over tuples in the user-permission relation extracted from the log, uses selected tuples as seeds for constructing candidate rules, and attempts to generalize each candidate rule to cover additional tuples in the user-permission relation by replacing conjuncts in attribute expressions with constraints. After constructing candidate rules that together cover the entire user-permission relation, it attempts to improve the policy by merging and simplifying candidate rules. Finally, it selects the highest-quality candidate rules for inclusion in the generated policy.

Several changes are needed to our algorithm for mining ABAC policies from ACLs to adapt it to mining from logs. When the algorithm generalizes, merges, or simplifies rules, it discards candidate rules that are invalid, i.e., that produce over-assignments. We modify those parts of the algorithm to consider those candidate rules, because, as discussed above, over-assignments must be permitted. To evaluate those candidate rules, we introduce generalized notions of rule quality and policy quality that quantify a trade-off between the number of over-assignments and other aspects of quality. We consider a metric that includes the normalized number of over-assignments in a weighted sum, a frequency-sensitive variant that assigns higher quality to rules that cover more frequently used entitlements, along the lines of [MPC12], and a metric based on a theory quality metric in inductive logic programming [MB00, MF01]. We also modify Xu and Stoller's approach to noise detection to take usage frequency into account.

ABAC policy mining is similar to inductive logic programming (ILP), which learns logic-programming rules from facts. Mining ABAC policies from logs and attribute data is similar to ILP algorithms for learning from positive examples, because those algorithms allow the learned rules to imply more than the given facts (i.e., in our terminology, to have over-assignments). We implemented a translation from ABAC policy mining to Progol [Mug97, MF01], a well-known ILP system.

We evaluated our algorithm and the ILP-based approach on some relatively small but non-trivial hand-written case studies and on synthetic ABAC policies. The results demonstrate our algorithm's effectiveness even when the log reflects only a fraction of the entitlements.

2 ABAC policy language

This section presents the ABAC policy language used in our work. It is adopted from [XS13a]. For simplicity and concreteness, we consider a specific ABAC policy language, but our approach is general and can be adapted to other ABAC policy languages. Our ABAC policy language contains all of the common ABAC policy language constructs, except arithmetic inequalities and negation. Extending our algorithm to handle those constructs is future work.

ABAC policies refer to attributes of users and resources. Given a set U of users and a set A_{u} of user attributes, user attribute data is represented by a function d_{u} such that $d_{\mathrm{u}}(u, a)$ is the value of attribute a for user u. There is a distinguished user attribute uid that has a unique value for each user. Similarly, given a set R of resources and a set A_{r} of resource attributes, resource attribute data is represented by a function d_{r} such that $d_{\mathrm{r}}(r, a)$ is the value of attribute a for resource r. There is a distinguished resource attribute rid that has a unique value for each resource. We assume the set A_{u} of user attributes can be partitioned into a set $A_{\mathrm{u},1}$ of *single-valued user attributes* which have atomic values, and a set $A_{\mathrm{u,m}}$ of *multi-valued user attributes* whose

values are sets of atomic values. Similarly, we assume the set A_{r} of resource attributes can be partitioned into a set $A_{\mathrm{r,1}}$ of *single-valued resource attributes* and a set of $A_{\mathrm{r,m}}$ of *multi-valued resource attributes*. Let $\mathrm{Val_s}$ be the set of possible atomic values of attributes. We assume $\mathrm{Val_s}$ includes a distinguished value $\perp$ used to indicate that an attribute's value is unknown. The set of possible values of multi-valued attributes is $\mathrm{Val_m} = \mathrm{Set}(\mathrm{Val_s} \setminus \{\perp\}) \cup \perp$, where $\mathrm{Set}(S)$ is the powerset of set S.

Attribute expressions are used to express the sets of users and resources to which a rule applies. A *user-attribute expression* (UAE) is a function e such that, for each user attribute a, $e(a)$ is either the special value $\top$, indicating that e imposes no constraint on the value of attribute a, or a set (interpreted as a disjunction) of possible values of a excluding $\perp$ (in other words, a subset of $\mathrm{Val_s} \setminus \{\perp\}$ or $\mathrm{Val_m} \setminus \{\perp\}$, depending on whether a is single-valued or multi-valued). We refer to the set $e(a)$ as the *conjunct* for attribute a. We say that expression e *uses* an attribute a if $e(a) \neq \top$. Let $\mathrm{attr}(e)$ denote the set of attributes used by e. Let $\mathrm{attr_1}(e)$ and $\mathrm{attr_m}(e)$ denote the sets of single-valued and multi-valued attributes, respectively, used by e.

A user u *satisfies* a user-attribute expression e, denoted $u \models e$, iff $(\forall a \in A_{\mathrm{u,1}}.\, e(a) = \top \vee \exists v \in e(a).\, d_{\mathrm{u}}(u, a) = v)$ and $(\forall a \in A_{\mathrm{u,m}}.\, e(a) = \top \vee \exists v \in e(a).\, d_{\mathrm{u}}(u, a) \supseteq v)$. For multi-valued attributes, we use the condition $d_{\mathrm{u}}(u, a) \supseteq v$ instead of $d_{\mathrm{u}}(u, a) = v$ because elements of a multi-valued user attribute typically represent some type of capabilities of a user, so using $\supseteq$ expresses that the user has the specified capabilities and possibly more.

For example, suppose $A_{\mathrm{u,1}} = \{\mathrm{dept}, \mathrm{position}\}$ and $A_{\mathrm{u,m}} = \{\mathrm{courses}\}$. The function e_1 with $e_1(\mathrm{dept}) = \{\mathrm{CS}\}$ and $e_1(\mathrm{position}) = \{\mathrm{grad}, \mathrm{ugrad}\}$ and $e_1(\mathrm{courses}) = \{\{\mathrm{CS101}, \mathrm{CS102}\}\}$ is a user-attribute expression satisfied by users in the CS department who are either graduate or undergraduate students and whose courses include CS101 and CS102.

We introduce a concrete syntax for attribute expressions, for improved readability in examples. We write a user attribute expression as a conjunction of the conjuncts not equal to $\top$. Suppose $e(a) \neq \top$. Let $v = e(a)$. When a is single-valued, we write the conjunct for a as $a \in v$; as syntactic sugar, if v is a singleton set $\{s\}$, we may write the conjunct as $a = s$. When a is multi-valued, we write the conjunct for a as $a \supseteq\!\in v$ (indicating that a is a superset of an element of v); as syntactic sugar, if v is a singleton set $\{s\}$, we may write the conjunct as $a \supseteq s$. For example, the above expression e_1 may be written as $\mathrm{dept} = \mathrm{CS} \wedge \mathrm{position} \in \{\mathrm{ugrad}, \mathrm{grad}\} \wedge \mathrm{courses} \supseteq \{\mathrm{CS101}, \mathrm{CS102}\}$. For an example that uses $\supseteq\!\in$, the expression e_2 that is the same as e_1 except with $e_2(\mathrm{courses}) = \{\{\mathrm{CS101}\}, \{\mathrm{CS102}\}\}$ may be written as $\mathrm{dept} = \mathrm{CS} \wedge \mathrm{position} \in \{\mathrm{ugrad}, \mathrm{grad}\} \wedge \mathrm{courses} \supseteq\!\in \{\{\mathrm{CS101}\}, \{\mathrm{CS102}\}\}$, and is satisfied by graduate or undergraduate students in the CS department whose courses include either CS101 or CS102.

The *meaning* of a user-attribute expression e, denoted $[\![e]\!]_U$, is the set of users in U that satisfy it: $[\![e]\!]_U = \{u \in U \mid u \models e\}$. User attribute data is an implicit argument to $[\![e]\!]_U$. We say that e *characterizes* the set $[\![e]\!]_U$.

A *resource-attribute expression* (RAE) is defined similarly, except using the set A_{r} of resource attributes instead of the set A_{u} of user attributes. The semantics of RAEs is defined similarly to the semantics of UAEs, except simply using equality, not $\supseteq$, in the condition for multi-valued attributes in the definition of "satisfies", because we do not interpret elements of multi-valued resource attributes specially (e.g., as capabilities).

Constraints are used to express relationships between users and resources. An *atomic constraint* is a formula f of the form $a_{\mathrm{u,m}} \supseteq a_{\mathrm{r,m}}$, $a_{\mathrm{u,m}} \ni a_{\mathrm{r,1}}$, or $a_{\mathrm{u,1}} = a_{\mathrm{r,1}}$, where $a_{\mathrm{u,1}} \in A_{\mathrm{u,1}}$, $a_{\mathrm{u,m}} \in A_{\mathrm{u,m}}$, $a_{\mathrm{r,1}} \in A_{\mathrm{r,1}}$, and $a_{\mathrm{r,m}} \in A_{\mathrm{r,m}}$. The first two forms express that user attributes contain specified values. This is a common type of constraint, because user attributes typically represent some type of capabilities of a user. Other forms of atomic constraint are possible (e.g., $a_{\mathrm{u,m}} \subseteq a_{\mathrm{r,m}}$) but less common, so we leave them for future work. Let $\mathrm{uAttr}(f)$ and $\mathrm{rAttr}(f)$ refer to the user attribute and resource attribute, respectively, used in f. User u and resource r *satisfy* an atomic constraint f, denoted $\langle u, r\rangle \models f$, if $d_{\mathrm{u}}(u, \mathrm{uAttr}(f)) \neq \perp$ and $d_{\mathrm{r}}(u, \mathrm{rAttr}(f)) \neq \perp$ and formula f holds when the values $d_{\mathrm{u}}(u, \mathrm{uAttr}(f))$ and $d_{\mathrm{r}}(u, \mathrm{rAttr}(f))$ are substituted in it.

A *constraint* is a set (interpreted as a conjunction) of atomic constraints. User u and resource r *satisfy* a constraint c, denoted $\langle u, r\rangle \models c$, if they satisfy every atomic constraint in c. In examples, we write constraints as conjunctions instead of sets. For example, the constraint "specialties $\supseteq$ topics $\wedge$ teams $\ni$ treatingTeam" is satisfied by user u and resource r if the user's specialties include all of the topics associated with the resource, and the set of teams associated with the user contains the treatingTeam associated with the resource.

A *user-permission tuple* is a tuple $\langle u, r, o\rangle$ containing a user, a resource, and an operation. This tuple means that user u has permission to perform operation o on resource r. A *user-permission relation* is a set of such tuples.

A *rule* is a tuple $\langle e_\mathrm{u}, e_\mathrm{r}, O, c\rangle$, where e_u is a user-attribute expression, e_r is a resource-attribute expression, O is a set of operations, and c is a constraint. For a rule $\rho = \langle e_\mathrm{u}, e_\mathrm{r}, O, c\rangle$, let $\mathrm{uae}(\rho) = e_\mathrm{u}$, $\mathrm{rae}(\rho) = e_\mathrm{r}$, $\mathrm{ops}(\rho) = O$, and $\mathrm{con}(\rho) = c$. For example, the rule ⟨true, type=task $\wedge$ proprietary=false, {read, request}, projects $\ni$ project $\wedge$ expertise $\supseteq$ expertise⟩ used in our project management case study can be interpreted as "A user working on a project can read and request to work on a non-proprietary task whose required areas of expertise are among his/her areas of expertise." User u, resource r, and operation o *satisfy* a rule ρ, denoted $\langle u, r, o\rangle \models \rho$, if $u \models \mathrm{uae}(\rho) \wedge r \models \mathrm{rae}(\rho) \wedge o \in \mathrm{ops}(\rho) \wedge \langle u, r\rangle \models \mathrm{con}(\rho)$.

An *ABAC policy* is a tuple $\langle U, R, Op, A_\mathrm{u}, A_\mathrm{r}, d_\mathrm{u}, d_\mathrm{r}, Rules\rangle$, where U, R, A_u, A_r, d_u, and d_r are as described above, Op is a set of operations, and $Rules$ is a set of rules.

The user-permission relation induced by a rule ρ is $[\![\rho]\!] = \{\langle u, r, o\rangle \in U \times R \times Op \mid \langle u, r, o\rangle \models \rho\}$. Note that U, R, d_u, and d_r are implicit arguments to $[\![\rho]\!]$.

The user-permission relation induced by a policy π with the above form is $[\![\pi]\!] = \bigcup_{\rho \in Rules} [\![\rho]\!]$.

3 Problem Definition

An *operation log entry* e is a tuple $\langle u, r, o, t\rangle$ where $u \in U$ is a user, $r \in R$ is a resource, $o \in Op$ is an operation, and t is a timestamp. An *operation log* is a sequence of operation log entries. The user-permission relation induced by an operation log L is $UP(L) = \{\langle u, r, o\rangle \mid \exists t.\ \langle u, r, o, t\rangle \in L\}$.

The input to the *ABAC-from-logs policy mining problem* is a tuple $I = \langle U, R, Op, A_\mathrm{u}, A_\mathrm{r}, d_\mathrm{u}, d_\mathrm{r}, L\rangle$, where U is a set of users, R is a set of resources, Op is a set of operations, A_u is a set of user attributes, A_r is a set of resource attributes, d_u is user attribute data, d_r is resource attribute data, and L is an operation log, such that the users, resources, and operations that appear in L are subsets of U, R, and Op, respectively. The goal of the problem is to find a set of rules $Rules$ such that the ABAC policy $\pi = \langle U, R, Op, A_\mathrm{u}, A_\mathrm{r}, d_\mathrm{u}, d_\mathrm{r}, Rules\rangle$ maximizes a suitable policy quality metric.

The policy quality metric should reflect the size and meaning of the policy. Size is measured by *weighted structural complexity* (WSC) [MCL+10], and smaller policies are considered to have higher quality. This is consistent with usability studies of access control rules, which conclude that more concise policies are more manageable [BM13]. Informally, the WSC of an ABAC policy is a weighted sum of the number of elements in the policy. Specifically, the WSC of an attribute expression is the number of atomic values that appear in it, the WSC of an operation set is the number of operations in it, the WSC of a constraint is the number of atomic constraints in it, and the WSC of a rule is a weighted sum of the WSCs of its components, namely, $\mathrm{WSC}(\langle e_\mathrm{u}, e_\mathrm{r}, O, c\rangle) = w_1\mathrm{WSC}(e_\mathrm{u}) + w_2\mathrm{WSC}(e_\mathrm{r}) + w_3\mathrm{WSC}(O) + w_4\mathrm{WSC}(c)$, where the w_i are user-specified weights. The WSC of a set of rules is the sum of the WSCs of its members.

The meaning $[\![\pi]\!]$ of the ABAC policy is taken into account by considering the differences from $UP(L)$, which consist of over-assignments and under-assignments. The over-assignments are $[\![\pi]\!] \setminus UP(L)$. The under-assignments are $UP(L) \setminus [\![\pi]\!]$. Since logs provide only a lower-bound on the actual user-permission relation (a.k.a entitlements), it is necessary to allow some over-assignments, but not too many. Allowing

under-assignments is beneficial if the logs might contain noise, in the form of log entries representing uses of permissions that should not be granted, because it reduces the amount of such noise that gets propagated into the policy, and it improves the stability of the generated policy. We define a policy quality metric that is a weighted sum of these aspects:

$$Q_{\text{pol}}(\pi, L) = \text{WSC}(\pi) + w_o \,|\, [\![\pi]\!] \setminus UP(L)| \,/\, |U| + w_u |UP(L) \setminus [\![\pi]\!] \,|_L \tag{1}$$

where the *policy over-assignment weight* w_o and the *under-assignment weight* w_u are user-specified weights for over-assignments and under-assignments, respectively, and for a set S of user-permission tuples, the frequency-weighted size of S with respect to log L is $|S|_L = \sum_{\langle u,r,o\rangle \in S} \text{freq}(\langle u, r, o\rangle, L)$, where the relative frequency of a user-permission tuple in a log is given by the *frequency function* $\text{freq}(\langle u, r, o\rangle, L) = |\{e \in L \mid \text{userPerm}(e) = \langle u, r, o\rangle\}| \,/\, |L|$, where the user-permission part of a log entry is given by $\text{userPerm}(\langle u, r, o, t\rangle) = \langle u, r, o\rangle$. Using the relative frequency, instead of the number of occurrences, of a user-permission tuple as the weight in the definition of frequency-weighted size allows w_u to be chosen independent of the duration of monitoring and the number of monitored users; for example, if the duration of monitoring or the number of monitored users is doubled, and the distribution of their behaviors and the value of w_u are kept unchanged, then the policy quality is unchanged. Similarly, dividing the number of over-assignments by $|U|$ allows w_o to be chosen independent of the number of users (and the duration of monitoring).

For simplicity, our presentation of the problem and algorithm assume that attribute data does not change during the time covered by the log. Accommodating changes to attribute data is not difficult. It mainly requires re-defining the notions of policy quality and rule quality (introduced in Section 4) to be based on the set of log entries covered by a rule, denoted $[\![\rho]\!]_{\text{LE}}$, rather than $[\![\rho]\!]$. The definition of $[\![\rho]\!]_{\text{LE}}$ is similar to the definition of $[\![\rho]\!]$, except that, when determining whether a log entry is in $[\![\rho]\!]_{\text{LE}}$, the attribute data in effect at the time of the log entry is used.

4 Algorithm

Our algorithm is based on the algorithm for mining ABAC policies from ACLs and attribute data in [XS13a]. Our algorithm does not take the order of log entries into account, so the log can be summarized by the user-permission relation UP_0 induced by the log and the frequency function freq, described in the penultimate paragraph of Section 3.

Top-level pseudocode appears in Figure 1. We refer to tuples selected in the first statement of the first while loop as *seeds*. The top-level pseudocode is explained by embedded comments. It calls several functions, described next. Function names hyperlink to pseudocode for the function, if it is included in the paper, otherwise to the description of the function.

The function addCandRule($s_{\text{u}}, s_r, s_{\text{o}}, cc, uncovUP, Rules$) in Figure 2 first calls computeUAE to compute a user-attribute expression e_{u} that characterizes s_{u}, and computeRAE to compute a resource-attribute expression e_{r} that characterizes s_r. It then calls generalizeRule($\rho, cc, uncovUP, Rules$) to generalize rule $\rho = \langle e_{\text{u}}, e_{\text{r}}, s_{\text{o}}, \emptyset\rangle$ to ρ' and adds ρ' to candidate rule set *Rules*. The details of the functions called by addCandRule are described next.

The function computeUAE(s, U) in Figure 3 computes a user-attribute expression e_{u} that characterizes the set s of users. Preference is given to attribute expressions that do not use uid, since attribute-based policies are generally preferable to identity-based policies, even when they have higher WSC, because attribute-based generalize better. After constructing a candidate expression e, it calls elimRedundantSets(e), which attempts to lower the WSC of e by examining the conjunct for each multi-valued user attribute, and removing each set that is a superset of another set in the same conjunct; this leaves the meaning of the rule unchanged, because $\supseteq$ is used in the condition for multi-valued attributes in the semantics of user attribute

expressions. Pseudocode for elimRedundantSets is straightforward and omitted. The expression e_{u} returned by computeUAE might not be minimum-sized among expressions that characterize s: it is possible that some attributes mapped to a set of values by e_{u} can instead be mapped to $\top$. We defer minimization of e_{u} until after the call to generalizeRule (described below), because minimizing e_{u} before that would reduce opportunities to find relations between values of user attributes and resource attributes in generalizeRule.

The function computeRAE(s, R) computes a resource-attribute expression that characterizes the set s of resources. The definition is the same as for computeUAE, except using resource attributes instead of user attributes, and the call to elimRedundantSets is omitted. Pseudocode for computeRAE is omitted.

The function candidateConstraint(r, u) returns a set containing all of the atomic constraints that hold between resource r and user u. Pseudocode for candidateConstraint is straightforward and omitted.

The function generalizeRule$(\rho, cc, \mathit{uncovUP}, \mathit{Rules})$ in Figure 4 attempts to generalize rule ρ by adding some of the atomic constraints in cc to ρ and eliminating the conjuncts of the user attribute expression and/or the resource attribute expression corresponding to the attributes used in those constraints, i.e., mapping those attributes to $\top$. We call a rule obtained in this way a *generalization* of ρ. Such a rule is more general than ρ in the sense that it refers to relationships instead of specific values. Also, the user-permission relation induced by a generalization of ρ is a superset of the user-permission relation induced by ρ. generalizeRule$(\rho, cc, \mathit{uncovUP}, \mathit{Rules})$ returns the generalization ρ' of ρ with the best quality according to a given rule quality metric. Note that ρ' may cover tuples that are already covered (i.e., are in UP); in other words, our algorithm can generate policies containing rules whose meanings overlap.

A *rule quality metric* is a function $Q_{\mathrm{rul}}(\rho, UP)$ that maps a rule ρ to a totally-ordered set, with the ordering chosen so that larger values indicate high quality. The second argument UP is a set of user-permission tuples. Our rule quality metric assigns higher quality to rules that cover more currently uncovered user-permission tuples and have smaller size, with an additional term that imposes a penalty for over-assignments, measured as a fraction of the number of user-permission tuples covered by the rule, and with a weight specified by a parameter w'_o, called the *rule over-assignment weight*.

$$Q_{\mathrm{rul}}(\rho, UP) = \frac{|\mathrm{covered}(\rho)|}{|\rho|} \times (1 - \frac{w'_o \times |\mathrm{overAssign}(\rho)|}{|\,[\![\rho]\!]\,|}).$$

In generalizeRule, *uncovUP* is the second argument to Q_{rul}, so $[\![\rho]\!] \cap UP$ is the set of user-permission tuples in UP_0 that are covered by ρ and not covered by rules already in the policy. The loop over i near the end of the pseudocode for generalizeRule considers all possibilities for the first atomic constraint in cc that gets added to the constraint of ρ. The function calls itself recursively to determine the subsequent atomic constraints in c that get added to the constraint.

We also developed a frequency-sensitive variant of this rule quality metric. Let $Q^{\mathrm{freq}}_{\mathrm{rul}}$ denote the frequency-weighted variant of Q_{rul}, obtained by weighting each user-permission tuple by its relative frequency (i.e., fraction of occurrences) in the log, similar to the definition of λ-distance in [MPC12]. Specifically, the definition of $Q^{\mathrm{freq}}_{\mathrm{rul}}$ is obtained from the definition of Q_{rul} by replacing $|\mathrm{covered}(\rho)|$ with $|\mathrm{covered}(\rho)|_L$ (recall that $|\cdot|_L$ is defined in Section 3).

We also developed a rule quality metric $Q^{\mathrm{ILP}}_{\mathrm{rul}}$ based closely on the theory quality metric for inductive logic programming described in [Mug95]. Details of the definition appear in Appendix A.

The function mergeRules(*Rules*) in Figure 4 attempts to improve the quality of *Rules* by removing redundant rules and merging pairs of rules. A rule ρ in *Rules* is *redundant* if *Rules* contains another rule ρ' such that every user-permission tuple in UP_0 that is in $[\![\rho]\!]$ is also in $[\![\rho']\!]$. Informally, rules ρ_1 and ρ_2 are merged by taking, for each attribute, the union of the conjuncts in ρ_1 and ρ_2 for that attribute. If adding the resulting rule ρ_{mrg} to the policy and removing rules (including ρ_1 and ρ_2) that become redundant improves policy quality and does not introduce over-assignments where none existed before, then ρ_{mrg} is

```
// Rules is the set of candidate rules
Rules = ∅
// uncovUP contains user-permission tuples
// in UP_0 that are not covered by Rules
uncovUP = UP_0.copy()
while ¬uncovUP.isEmpty()
  // Select an uncovered tuple as a "seed".
  ⟨u, r, o⟩ = some tuple in uncovUP
  cc = candidateConstraint(r, u)
  // s_u contains users with permission ⟨r, o⟩
  // and that have the same candidate
  // constraint for r as u
  s_u = {u' ∈ U | ⟨u', r, o⟩ ∈ UP_0
            ∧ candidateConstraint(r, u') = cc}
  addCandRule(s_u, {r}, {o}, cc, uncovUP, Rules)
  // s_o is set of operations that u can apply to r
  s_o = {o' ∈ Op | ⟨u, r, o'⟩ ∈ UP_0}
  addCandRule({u}, {r}, s_o, cc, uncovUP, Rules)
end while
```

```
// Repeatedly merge and simplify
// rules, until this has no effect
mergeRules(Rules)
while simplifyRules(Rules)
      && mergeRules(Rules)
  skip
end while
// Select high quality rules into Rules'.
Rules' = ∅
Repeatedly move highest-quality rule
from Rules to Rules' until
∑_{ρ∈Rules'} [[ρ]] ⊇ UP_0, using
UP_0 \ [[Rules']] as second argument to
Q_rul, and discarding a rule if it does
not cover any tuples in UP_0 currently
uncovered by Rules'.
return Rules'
```

Figure 1: Policy mining algorithm. The pseudocode starts in column 1 and continues in column 2.

```
function addCandRule(s_u, s_r, s_o, cc, uncovUP, Rules)
// Construct a rule ρ that covers user-permission tuples {⟨u, r, o⟩ | u ∈ s_u ∧ r ∈ s_r ∧ o ∈ s_o}.
e_u = computeUAE(s_u, U);
e_r = computeRAE(s_r, R);
ρ = ⟨e_u, e_r, s_o, ∅⟩
ρ' = generalizeRule(ρ, cc, uncovUP, Rules);
Rules.add(ρ');
uncovUP.removeAll([[ρ']])
```

Figure 2: Compute a candidate rule ρ' and add ρ' to candidate rule set *Rules*

added to *Rules*, and the redundant rules are removed from *Rules*. As optimizations (in the implementation, not reflected in the pseudocode), meanings of rules are cached, and policy quality is computed incrementally. mergeRules(*Rules*) updates its argument *Rules* in place, and it returns a Boolean indicating whether any rules were merged.

5 Functions to Simplify Rules

The function simplifyRules(*Rules*) in Figure 5 attempts to simplify all of the rules in *Rules*. It updates its argument *Rules* in place, replacing rules in *Rules* with simplified versions when simplification succeeds. It returns a Boolean indicating whether any rules were simplified. It attempts to simplify each rule in several ways, which are embodied in the following functions that it calls. The names of these functions start with "elim", because they attempt to eliminate unnecessary parts of rules. To enable simplifyRules to determine whether any rules were simplified, each "elim" function returns a Boolean value indicating whether it simplified any rules. For brevity, computation of the Boolean return values of "elim" functions are not reflected in the pseudocode.

The function elimRedundantSets is described above. It returns false, even if some redundant sets were

function computeUAE(s, U)
// *First try to characterize set s of users without using*
// uid. *Use all other attributes which have known*
// *values for all users in s.*
$e = (\lambda\, a \in A_{\mathrm{u}}.$
  $a = \mathrm{uid} \vee (\exists\, u \in s.\, d_{\mathrm{u}}(u, a) = \bot)\, ?\, \top : \bigcup_{u \in s} d_{\mathrm{u}}(u, a))$
if $[\![e]\!]_U \neq s$
  // uid *is needed to characterize s*
  $e(\mathrm{uid}) = \bigcup_{u \in s} d_{\mathrm{u}}(u, \mathrm{uid})$
end if
elimRedundantSets(e)
return e

Figure 3: Compute a user-attribute expression that characterizes set s of users, where U is the set of all users.

eliminated, because elimination of redundant sets does not affect the meaning or mergeability of rules, so it need not trigger another iteration of merging and simplification.

The function elimConjuncts(ρ, *Rules*, *UP*) in Figure 5 attempts to increase the quality of rule ρ by eliminating some of the conjuncts. It calls the function elimConjunctsHelper(ρ, A, *UP*) in Figure 5, which considers all rules that differ from ρ by mapping a subset A' of the tagged attributes in A to $\top$ instead of to a set of values; among the resulting rules that are valid, it returns one with the highest quality. A *tagged attribute* is a pair of the form $\langle$"user", $a\rangle$ with $a \in A_{\mathrm{u}}$ or $\langle$"res", $a\rangle$ with $a \in A_{\mathrm{r}}$. The set A_{unrm} in elimConjuncts is a set of *unremovable* tagged attributes; it is a parameter of the algorithm, specifying attributes that should not be eliminated, because eliminating them increases the risk of generating an overly general policy, i.e., a policy that might grant inappropriate permissions when new users or new resources (hence new permissions) are added to the system. We use a combinatorial algorithm for elimConjuncts that evaluates all combinations of conjuncts that can be eliminated, because elimination of one conjunct might prevent elimination of another conjunct. This algorithm makes elimConjuncts worst-case exponential in the numbers of user attributes and resource attributes that can be eliminated while preserving validity of the rule; in practice the number of such attributes is small. elimConjuncts also considers whether to remove conjuncts from the user attribute expression or the resource attribute expression first, because elimination of a conjunct in one attribute expression might prevent elimination of a conjunct in the other. The algorithm could simply try both orders, but instead it uses a heuristic that, in our experiments, is faster and almost as effective: if maxConjunctSz(e_{u}) $\geq$ maxConjunctSz(e_{r}) then eliminate conjuncts from the user attribute expression first, otherwise eliminate conjuncts from the resource attribute expression first, where maxConjunctSz(e) is the size (WSC) of the largest conjunct in attribute expression e. elimConjunctsHelper calls the function elimAttribute(ρ, ta), which returns a copy of rule ρ with the conjunct for tagged attribute ta removed from the user attribute expression or resource attribute expression as appropriate (in other words, the specified attribute is mapped to $\bot$); pseudocode for elimAttribute is straightforward and omitted.

The function elimConstraints(ρ, *Rules*, *UP*) attempts to improve the quality of ρ by removing unnecessary atomic constraints from ρ's constraint. An atomic constraint is *unnecessary* in a rule ρ if removing it from ρ's constraint leaves ρ valid. Pseudocode for elimConstraints is analogous to elimConjuncts, except it considers removing atomic constraints instead of conjuncts from rules.

The function elimElements(ρ) attempts to decrease the WSC of rule ρ by removing elements from sets in conjuncts for multi-valued user attributes, if removal of those elements preserves validity of ρ; note that, because $\subseteq$ is used in the semantics of user attribute expressions, the set of user-permission tuples that satisfy a rule is never decreased by such removals. It would be reasonable to use a combinatorial algorithm

function generalizeRule(ρ, *cc*, *uncovUP*, *Rules*)
// ρ_{best} *is best generalization of* ρ
$\rho_{\text{best}} = \rho$
// *gen*[*i*][*j*] *is a generalization of* ρ *using*
// *cc*′[*i*]
gen = new Rule[*cc*.length][3]
for $i = 1$ **to** *cc*.length
  $f = cc[i]$
  // *generalize by adding* f *and eliminating*
  // *conjuncts for both attributes used in* f.
  $gen[i][1] = \langle \text{uae}(\rho)[\text{uAttr}(f) \mapsto \top],$
    $\text{rae}(\rho)[\text{rAttr}(f) \mapsto \top],$
    $\text{ops}(\rho), \text{con}(\rho) \cup \{f\}\rangle$
  // *generalize by adding* f *and eliminating*
  // *conjunct for user attribute used in* f
  $gen[i][2] = \langle \text{uae}(\rho)[\text{uAttr}(f) \mapsto \top], \text{rae}(\rho),$
    $\text{ops}(\rho), \text{con}(\rho) \cup \{f\}\rangle$
  // *generalize by adding* f *and eliminating*
  // *conjunct for resource attrib. used in* f.
  $gen[i][3] = \langle \text{uae}(\rho), \text{rae}(\rho)[\text{rAttr}(f) \mapsto \top],$
    $\text{ops}(\rho), \text{con}(\rho) \cup \{f\}\rangle$
end for
for $i = 1$ **to** *cc*.length and $j = 1$ **to** 3
  // *try to further generalize* *gen*[*i*]
  $\rho'' = \text{generalizeRule}(gen[i][j], cc[i{+}1\,..], uncovUP, Rules)$
  if $Q_{\text{rul}}(\rho'', uncovUP) > Q_{\text{rul}}(\rho_{\text{best}}, uncovUP)$
    $\rho_{\text{best}} = \rho''$
  end if
end for
return ρ_{best}

function mergeRules(*Rules*)
// *Remove redundant rules*
$redun = \{\rho \in Rules \mid \exists\, \rho' \in Rules \setminus \{\rho\}.$
  $[\![\rho]\!] \cap UP_0 \subseteq [\![\rho']\!] \cap UP_0\}$
Rules.removeAll(*redun*)
// *Merge rules*
$workSet = \{(\rho_1, \rho_2) \mid \rho_1 \in Rules \wedge \rho_2 \in Rules$
  $\wedge\, \rho_1 \neq \rho_2 \wedge \text{con}(\rho_1) = \text{con}(\rho_2)\}$
while not(*workSet*.empty())
  $(\rho_1, \rho_2) = workSet.\text{remove}()$
  $\rho_{\text{mrg}} = \langle \text{uae}(\rho_1) \cup \text{uae}(\rho_2),$
    $\text{rae}(\rho_1) \cup \text{rae}(\rho_2),$
    $\text{ops}(\rho_1) \cup \text{ops}(\rho_2), \text{con}(\rho_1)\rangle$
  // *Find rules that become redundant*
  // *if merged rule* ρ_{mrg} *is added*
  $redun = \{\rho \in Rules \mid [\![\rho]\!] \subseteq [\![\rho_{\text{mrg}}]\!]\}$
  // *Add the merged rule and remove redun-*
  // *dant rules if this improves policy quality*
  // *and does not introduce over-assignments.*
  // *where none existed before.*
  if $Q_{\text{pol}}(Rules \cup \{\rho_{\text{mrg}}\} \setminus redun) < Q_{\text{pol}}(Rules)$
    $\wedge\, (\text{noOA}(\rho_1) \wedge \text{noOA}(\rho_2) \Rightarrow \text{noOA}(\rho_{\text{mrg}}))$
    Rules.removeAll(*redun*)
    workSet.removeAll($\{(\rho_1, \rho_2) \in workSet \mid$
      $\rho_1 \in redun \vee \rho_2 \in redun\}$)
    workSet.addAll($\{(\rho_{\text{mrg}}, \rho) \mid \rho \in Rules$
      $\wedge\, \text{con}(\rho) = \text{con}(\rho_{\text{mrg}})\}$)
    Rules.add(ρ_{mrg})
  end if
end while
return true if any rules were merged

Figure 4: Left: Generalize rule ρ by adding some formulas from *cc* to its constraint and eliminating conjuncts for attributes used in those formulas. $f[x \mapsto y]$ denotes a copy of function f modified so that $f(x) = y$. $a[i..]$ denotes the suffix of array a starting at index i. Right: Merge pairs of rules in *Rules*, when possible, to reduce the WSC of *Rules*. (a, b) denotes an unordered pair with components a and b. The union $e = e_1 \cup e_2$ of attribute expressions e_1 and e_2 over the same set A of attributes is defined by: for all attributes a in A, if $e_1(a) = \top$ or $e_2(a) = \top$ then $e(a) = \top$ otherwise $e(a) = e_1(a) \cup e_2(a)$. $\text{noOA}(\rho)$ holds if ρ has no over-assignments, i.e., $[\![\rho]\!] \subseteq UP_0$.

for elimElements, in the same style as elimConjuncts and elimConstraints, because elimination of one set element can prevent elimination of another. We decided to use a simple linear algorithm for this function, for simplicity and because it is likely to give the same results, because elimElements usually eliminates only 0 or 1 set elements per rule in our experiments. Pseudocode for elimElements is straightforward and omitted.

The function elimOverlapVal(ρ, *Rules*) attempts to decrease the WSC of rule ρ by removing values from conjuncts of attribute expressions in ρ if there are other rules that cover the affected user-permission tuples and have higher quality. Specifically, a value v in the conjunct for a user attribute a in ρ is removed if there is another rule ρ' in *Rules* such that (1) $\text{attr}(\text{uae}(\rho')) \subseteq \text{attr}(\text{uae}(\rho))$ and $\text{attr}(\text{rae}(\rho')) \subseteq \text{attr}(\text{rae}(\rho))$, (2)

the conjunct of uae(ρ') for a contains v, (3) each conjunct of uae(ρ') or rae(ρ') other than the conjunct for a is either $\top$ or a superset of the corresponding conjunct of ρ, and (4) con(ρ') $\subseteq$ con(ρ). The condition for removal of a value in the conjunct for a resource attribute is analogous. If a conjunct of uae(ρ) or rae(ρ) becomes empty, ρ is removed from *Rules*. elimOverlapVal(ρ, *Rules*) returns true if it modifies or removes ρ, otherwise it returns false. Pseudocode for elimOverlapVal is straightforward and omitted.

The function elimOverlapOp(ρ, *Rules*) attempts to decrease the WSC of rule ρ by removing operations from ops(ρ), if there are other rules that cover the affected user-permission tuples. Specifically, an operation o is removed from ops(ρ) if there is another rule ρ' in *Rules* such that (1) attr(uae(ρ')) $\subseteq$ attr(uae(ρ)) and attr(rae(ρ')) $\subseteq$ attr(rae(ρ)), (2) ops(ρ') contains o, (3) each conjunct of uae(ρ') or rae(ρ') is either $\top$ or a superset of the corresponding conjunct of ρ, and (4) con(ρ') is a subset of con(ρ). If ops(ρ) becomes empty, ρ is removed from *Rules*. elimOverlapOp(ρ, *Rules*) returns true if it modifies or removes ρ, otherwise it returns false. Pseudocode for elimOverlapOp is straightforward and omitted.

Noise Detection. If the input might contain noise, in the form of log entries representing uses of permissions that should not be granted, then the algorithm should allow the generated policy to have under-assignments (i.e., it should allow $UP(L) \setminus [\![\pi]\!]$ to be non-empty), to reduce the amount of such noise that gets propagated into the policy. Such noise in the logs is due to over-assignments in the currently deployed access control policy. We adopt the approach in [XS13a] to identify suspected over-assignment noise. The main idea is that rules with quality below a threshold τ are considered to represent noise and hence are omitted from the generated policy, and that user-permission tuples in $UP(L)$ that are represented only by rules classified as noise are reported as suspected over-assignment noise. An additional consideration in this context is that usage frequency should be considered: under-assignments in the generated policy should correspond to infrequently used entitlements. Thus, the rule quality metric used to identify rules that represent noise should be frequency-sensitive. One option is to use the frequency-sensitive rule quality metric $Q_{\text{rul}}^{\text{freq}}$. Another option is to use a metric focused exclusively on the usage frequency of the permissions covered by a rule, ignoring the size of the rule, the number of covered user-permission tuples, and the over-assignments (if any). Such a metric is $Q_{\text{freq}}(\rho, L) = |\,[\![\rho]\!]\,|^{-1} \sum_{t \in [\![\rho]\!]} \text{freq}(t, L)$.

The rule quality metric used for noise detection can be the same rule quality metric used in the main part of the algorithm, in which case the main part of the algorithm is modified to stop adding candidate rules to the final policy when the rule quality drops below threshold, or a different metric, in which case an additional loop is added at the end of the algorithm to find and remove rules whose quality is below threshold.

6 Evaluation Methodology

We evaluate our policy mining algorithms on synthetic operation logs generated from an ABAC policy (some handwritten and some synthetic) and probability distributions characterizing the frequency of actions. This allows us to evaluate the effectiveness of our algorithm by comparing the mined policies with the original ABAC policies. We are eager to also evaluate our algorithm on actual operation logs, but unfortunately, we have not yet been able to obtain access to actual operation logs for any system with a non-trivial ABAC policy and a significant number of users.

6.1 ABAC Policies

Case Studies. We developed four case studies for use in evaluation of our algorithm, described briefly here. Details of the case studies, including all policy rules, various size metrics (number of users, number of resources, etc.), and some illustrative attribute data, appear in [XS13a].

function simplifyRules(*Rules*)
for ρ **in** *Rules*
  elimRedundantSets(uae(ρ))
  elimConjuncts(ρ, *Rules*, UP_0)
  elimElements(ρ)
end for
for ρ **in** *Rules*
  elimOverlapVal(ρ, *Rules*)
end for
for ρ **in** *Rules*
  elimOverlapOp(ρ, *Rules*)
end for
for ρ **in** *Rules*
  elimConstraints(ρ, *Rules*, UP_0)
end for
return true if any ρ in *Rules* was changed

function elimConjuncts(ρ, *Rules*, *UP*)
$Au = \{\texttt{"user"}\} \times \text{attr}(\text{uae}(\rho)) \setminus A_{\text{unrm}}$
$Ar = \{\texttt{"res"}\} \times \text{attr}(\text{rae}(\rho)) \setminus A_{\text{unrm}}$
if maxConjunctSz(uae(ρ)) $\geq$ maxConjunctSz(rae(ρ))
  $\rho' =$ elimConjunctsHelper(ρ, *Au*, *UP*)
  $\rho'' =$ elimConjunctsHelper(ρ', *Ar*, *UP*)
else
  $\rho' =$ elimConjunctsHelper(ρ, *Ar*, *UP*)
  $\rho'' =$ elimConjunctsHelper(ρ', *Au*, *UP*)
end if
if $\rho'' \neq \rho$
  replace ρ with ρ'' in *Rules*
end if

function elimConjunctsHelper(ρ, *A*, *UP*)
$\rho_{\text{best}} = \rho$
// Discard tagged attributes ta such that elimi-
// nation of the conjunct for ta makes ρ invalid.
for *ta* **in** *A*
  $\rho' =$ elimAttribute(ρ, *ta*)
  if not $[\![\rho']\!] \subseteq UP_0$
    A.remove(*ta*)
  end if
end for
for $i = 1$ **to** *A*.length *// treat A as an array*
  $\rho' =$ elimAttribute(ρ, $A[i]$)
  $\rho'' =$ elimConjunctsHelper(ρ', $A[i{+}1\,..]$)
  if $Q_{\text{rul}}(\rho'', UP) > Q_{\text{rul}}(\rho_{\text{best}}, UP)$
    $\rho_{\text{best}} = \rho''$
  end if
end for
return ρ_{best}

Figure 5: Functions used to simplify rules.

Our *university case study* is a policy that controls access by students, instructors, teaching assistants, registrar officers, department chairs, and admissions officers to applications (for admission), gradebooks, transcripts, and course schedules. Our *health care case study* is a policy that controls access by nurses, doctors, patients, and agents (e.g., a patient's spouse) to electronic health records (HRs) and HR items (i.e., entries in health records). Our *project management case study* is a policy that controls access by department managers, project leaders, employees, contractors, auditors, accountants, and planners to budgets, schedules, and tasks associated with projects. Our *online video case study* is a policy that controls access to videos by users of an online video service.

The number of rules in the case studies is relatively small (10 ± 1 for the first three case studies, and 6 for online video), but they express non-trivial policies and exercise all the features of our policy language, including use of set membership and superset relations in attribute expressions and constraints. The manually written attribute dataset for each case study contains a small number of instances of each type of user and resource.

For the first three case studies, we generated a series of synthetic attribute datasets, parameterized by a number N, which is the number of departments for the university and project management case studies, and the number of wards for the health care case study. The generated attribute data for users and resources associated with each department or ward are similar to but more numerous than the attribute data in the manually written datasets. We did not bother creating synthetic data for the online video case study, because the rules are simpler.

Synthetic Policies. We generated synthetic policies using the algorithm proposed by Xu and Stoller [XS13a]. Briefly, the policy synthesis algorithm first generates the rules and then uses the rules to guide generation of the attribute data; this allows control of the number of granted permissions. The algorithm takes N_{rule}, the desired number of rules, as an input. The numbers of users and resources are not specified directly but are proportional to the number of rules. Generation of the rules and attribute data is based on several statistical distributions, which are either based loosely on the case studies or assumed to have a simple functional form (e.g., uniform distribution).

6.2 Log Generation

We first describe a straightforward but inefficient log generation algorithm, and then describe the log summary generation algorithm that is actually used in our experiments.

Log Generation Algorithm. The inputs to the algorithm are an ABAC policy π, the desired completeness of the log, and several probability distributions. The *completeness* of a log, relative to an ABAC policy, is the fraction of user-permission tuples in the meaning of the policy that appear in at least one entry in the log. We use completeness as a measure of log size relative to policy size. The algorithm generates each log entry by first selecting an ABAC rule, according to a probability distribution on roles, and then selecting a user-permission tuple that satisfies the rule, according to probability distributions on users, resources, and operations. This process is repeated until the specified completeness is reached.

Specifically, our log generation algorithm uses the following probability distributions. For each rule ρ in π, $P_{\text{rule}}(\rho)$ is the probability of selecting ρ. For each operation o and rule ρ, $P_{\text{op}}(o|\rho)$ is the conditional probability of selecting operation o when ρ has been selected; for sanity, we require $P_{\text{op}}(o|\rho) = 0$ if $o \notin \text{ops}(\rho)$. For each user u, resource r, and rule ρ, $P_{\text{ur}}(u, r|\rho)$ is the conditional probability of selecting user u and resource r when rule ρ and resource r have been selected; for sanity, we require $P_{\text{ur}}(u, r|\rho) = 0$ if $u \not\models \text{uae}(\rho) \vee r \not\models \text{rae}(\rho) \vee \langle u, r\rangle \not\models \text{con}(\rho)$. The algorithm generates each log entry as follows: select a rule ρ based on $P_{\text{rule}}(\rho)$, select an operation o based on $P_{\text{op}}(o|\rho)$ and a user-permission tuple $\langle u, r\rangle$ based on

$P_{\text{ur}}(u, r|\rho)$, and then create a log entry $\langle u, r, o, i\rangle$, where the unique identifier i is the index of this log entry in the synthetic log.

For convenience, we allow the user to specify probability distributions $P_{\text{user}}(u)$, $P_{\text{res}}(r)$, and $P'_{\text{op}}(o)$ instead of P_{ur} and P_{op}, and we compute P_{ur} and P_{op} from them as described below; intuitively, the probability associated with invalid values is distributed uniformly among the valid values by this computation. These alternative probability distributions provide less detailed control but are simpler and more convenient.

$$
\begin{aligned}
P_{\text{ur}}(u, r|\rho) &= \begin{cases} c^{-1} P_{\text{user}}(u) P_{\text{res}}(r) & \text{if } \langle u, r\rangle \in [\![\rho]\!]_{\text{ur}} \\ 0 & \text{otherwise} \end{cases} \\
&\quad \text{where } c = \sum_{\langle u,r\rangle \in [\![\rho]\!]_{\text{ur}}} P_{\text{user}}(u) P_{\text{res}}(r) \\
P_{\text{op}}(o|\rho) &= \begin{cases} c^{-1} P'_{\text{op}}(o) & \text{if } o \in \text{ops}(\rho) \\ 0 & \text{otherwise} \end{cases} \\
&\quad \text{where } c = \sum_{o \in \text{ops}(\rho)} P'_{\text{op}}(o)
\end{aligned}
$$

where the user-resource part of the meaning of a rule is $[\![\rho]\!]_{\text{ur}} = \bigcup_{\langle u,r,o\rangle \in [\![\rho]\!]} \{\langle u, r\rangle\}$.

Log Summary Generation Algorithm. When the probability distributions give much higher probability to some user-permission tuples than others, using the above algorithm to generate logs with high completeness is very inefficient, because very long logs, with many entries for high-probability user-permission tuples, are generated. Therefore, we adopt a different approach that takes advantage of the fact that our policy mining algorithm is insensitive to the order of log entries and depends only on the frequency of each user-permission tuple in the log. Specifically, instead of generating logs, we directly generate *log summaries*, which are simply sets of user-permission tuples with associated frequencies that sum to 1 (in other words, a log summary defines a set of user-permission tuples and a frequency function). For a given ABAC policy and probability distributions, we generate a log summary with completeness 1 using a program that computes the asymptotic frequency with which each tuple in UP_0 would appear in an infinitely long log generated using the above algorithm. To generate a series of log summaries with varying completeness, we start with the log summary with completeness 1 and repeatedly apply the following procedure that generates a log summary of lower completeness from a log summary with higher completeness: select the appropriate number of user-permission tuples from the latter, with the selection probability for each tuple equal to its associated frequency, and then normalize the frequencies of the selected tuples so they sum to 1.

Probability Distributions. An important characteristic of the probability distributions used in synthetic log and log summary generation is the ratio between the most frequent (i.e., most likely) and least frequent items of each type (rule, user, etc.). For case studies with manually written attribute data, we manually created probability distributions in which this ratio ranges from about 3 to 6. For case studies with synthetic data and synthetic policies, we generated probability distributions in which this ratio is 25 for rules, 25 for resources, 3 for users, and 3 for operations (the ratio for operations has little impact, because it is relevant only when multiple operations appear in the same rule, which is uncommon).

6.3 Metrics

For each case study and each associated attribute dataset (manually written or synthetic), we generate a synthetic operation log using the algorithm in Section 6.2 and then run our ABAC policy mining algorithms.

We evaluate the effectiveness of each algorithm by comparing the generated ABAC policy to the original ABAC policy, using the metrics described below.

Syntactic Similarity. The Jaccard similarity of two sets is $J(S_1, S_2) = |S_1 \cap S_2| \,/\, |S_1 \cup S_2|$. The syntactic similarity of UAEs is defined by $S^{\mathrm{u}}_{\mathrm{syn}}(e, e') = |A_{\mathrm{u}}|^{-1} \sum_{a \in A_{\mathrm{u}}} J(e(a), e'(a))$. The syntactic similarity of RAEs is defined by $S^{\mathrm{r}}_{\mathrm{syn}}(e, e') = |A_{\mathrm{r}}|^{-1} \sum_{a \in A_{\mathrm{r}}} J(e(a), e'(a))$. The syntactic similarity of rules $\langle e_{\mathrm{u}}, e_{\mathrm{r}}, O, c\rangle$ and $\langle e'_{\mathrm{u}}, e'_{\mathrm{r}}, O', c'\rangle$ is the average of the similarities of their components, specifically, the average of $S^{\mathrm{u}}_{\mathrm{syn}}(e_{\mathrm{u}}, e'_{\mathrm{u}})$, $S^{\mathrm{r}}_{\mathrm{syn}}(e_{\mathrm{r}}, e'_{\mathrm{r}})$, $J(O, O')$, and $J(c, c')$. The *syntactic similarity* of rule sets *Rules* and *Rules*$'$ is the average, over rules ρ in *Rules*, of the syntactic similarity between ρ and the most similar rule in *Rules*$'$. The *syntactic similarity* of policies π and π' is the maximum of the syntactic similarities of the sets of rules in the policies, considered in both orders (this makes the relation symmetric). Formally,

$$\begin{aligned} S_{\mathrm{syn}}(\mathit{Rules}, \mathit{Rules}') &= |\mathit{Rules}|^{-1} \sum_{\rho \in \mathit{Rules}} \min(\{S_{\mathrm{syn}}(\rho, \rho') \mid \rho' \in \mathit{Rules}'\})) \\ S_{\mathrm{syn}}(\pi, \pi') &= \max(S_{\mathrm{syn}}(\mathrm{rules}(\pi), \mathrm{rules}(\pi')), S_{\mathrm{syn}}(\mathrm{rules}(\pi'), \mathrm{rules}(\pi))) \end{aligned}$$

Syntactic similarity ranges from 0 (completely different) to 1 (identical).

Semantic Similarity. Semantic similarity measures the similarity of the entitlements granted by two policies. The *semantic similarity* of policies π and π' is defined by $J(\llbracket \pi \rrbracket, \llbracket \pi' \rrbracket)$. Semantic similarity ranges from 0 (completely different) to 1 (identical).

Fractions of Under-Assignments and Over-Assignments. To characterize the semantic differences between an original ABAC policy π_0 and a mined policy π in a way that distinguishes under-assignments and over-assignments, we compute the fraction of over-assignments and the fraction of under-assignments, defined by $|\llbracket \pi \rrbracket \setminus \llbracket \pi_0 \rrbracket| \,/\, |\llbracket \pi \rrbracket|$ and $|\llbracket \pi_0 \rrbracket \setminus \llbracket \pi \rrbracket| \,/\, |\llbracket \pi \rrbracket|$, respectively.

6.4 Noise

We consider two kinds of noise: permission noise in form of over-assignments (i.e., incorrectly granted permissions), and attribute noise. We do not consider under-assignments (i.e., absence of permissions that should be granted), because introducing under-assignments is equivalent to decreasing the log completeness.

To evaluate the effectiveness of our noise detection techniques in the presence of over-assignment noise with level ν_o and attribute noise with level ν_a, we start with an ABAC policy, change attribute values until the validity of $\nu_a |UP_0|$ user-permission tuples is changed, generate a synthetic log summary from the given ABAC rules and the modified attribute data, with the log summary generation algorithm modified to reflect that each log entry has a probability ν_o of embodying an over-assignment instead of a valid user-permission tuple, and then apply our policy mining algorithm to the resulting log summary and the modified attribute data. The over-assignments are invalid user-permission tuples generated according to the same probability distributions used in synthetic log generation. The changed attribute values are divided equally between missing values (i.e., replace a non-bottom value with bottom) and incorrect values (i.e., replace a non-bottom value with another non-bottom value).

Our current noise detection technique does not attempt to distinguish over-assignment noise from attribute noise (this is a topic for future research); policy analysts are responsible for determining whether a reported suspected error is due to an incorrect permission, an incorrect or missing attribute value, or a false alarm. Therefore, when comparing actual over-assignments to reported suspected over-assignments,

permission changes due to attribute noise (i.e., changes in the set of user-permission tuples that satisfy the policy rules) are included in the actual noise.

7 Experimental Results

This section presents experimental results using an implementation of our algorithm in Java. The implementation, case studies, and synthetic policies used in the experiments are available at `http://www.cs.stonybrook.edu/~stoller/`.

Over-Assignment Weight. The optimal choice for the over-assignment weights w_o and w'_o in the policy quality and rule quality metrics, respectively, depends on the log completeness. When log completeness is higher, fewer over-assignments are desired, and larger over-assignments weights give better results. In experiments, we take $w_o = c/2 - 15$ and $w'_o = w_o/10$, where c is log completeness. In a production setting, the exact log completeness would be unknown, but a rough estimate suffices, because our algorithm's results are robust to error in this estimate. For example, for case studies with manually written attribute data, when the actual log completeness is 80%, and the estimated completeness used to compute w_o varies from 70% to 90%, the semantic similarity of the original and mined policies varies by 0.04, 0.02, and 0 for university, healthcare, and project management, respectively.

Experimental Results. Figure 6 shows results from our algorithm. In each graph, curves are shown for the university, healthcare, and project management case studies with synthetic attribute data with N equal to 6, 10, and 10, respectively (average over results for 10 synthetic datasets, with 1 synthetic log per synthetic dataset), the online video case study with manually written attribute data (average over results for 10 synthetic logs), and synthetic policies with $N_{\text{rule}} = 20$ (average over results for 10 synthetic policies, with 1 synthetic log per policy). Error bars show standard deviation. Running time is at most 12 sec for each problem instance in our experiments.

For log completeness 100%, all four case study policies are reconstructed exactly, and the semantics of synthetic policies is reconstructed almost exactly: the semantic similarity is 0.98. This is a non-trivial result, especially for the case studies: an algorithm could easily generate a policy with over-assignments or generate more complex rules. As expected, the results get worse as log completeness decreases. When evaluating the results, it is important to consider what levels of log completeness are likely to be encountered in practice. One datapoint comes from Molloy *et al.*'s work on role mining from real logs [MPC12]. For the experiments in [MPC12, Tables 4 and 6], the actual policy is not known, but their algorithm produces policies with 0.52% or fewer over-assignments relative to $UP(L)$, and they interpret this as a good result, suggesting that they consider the log completeness to be near 99%. Based on this, we consider our experiments with log completeness below 90% to be severe stress tests, and results for log completeness 90% and higher to be more representative of typical results in practice.

Syntactic similarity for all four case studies is above 0.87 for log completeness 60% or higher, and is above 0.93 for log completeness 80% or higher. Syntactic similarity is lower for synthetic policies, but this is actually a good result. The synthetic policies tend to be unnecessarily complicated, and the mined policies are better in the sense that they have lower WSC. For example, for 100% log completeness, the mined policies have 0.98 semantic similarity to the synthetic policies (i.e., the meaning is almost the same), but the mined policies are simpler, with WSC 17% less than the original synthetic policies.

Semantic similarity is above 0.7 for log completeness 60% or higher, and above 0.89 for log completeness 80% or higher, for synthetic policies and for case studies other than healthcare. The semantic similarity is lower for healthcare, because the over-assignment weight given by the above formula is not optimal for

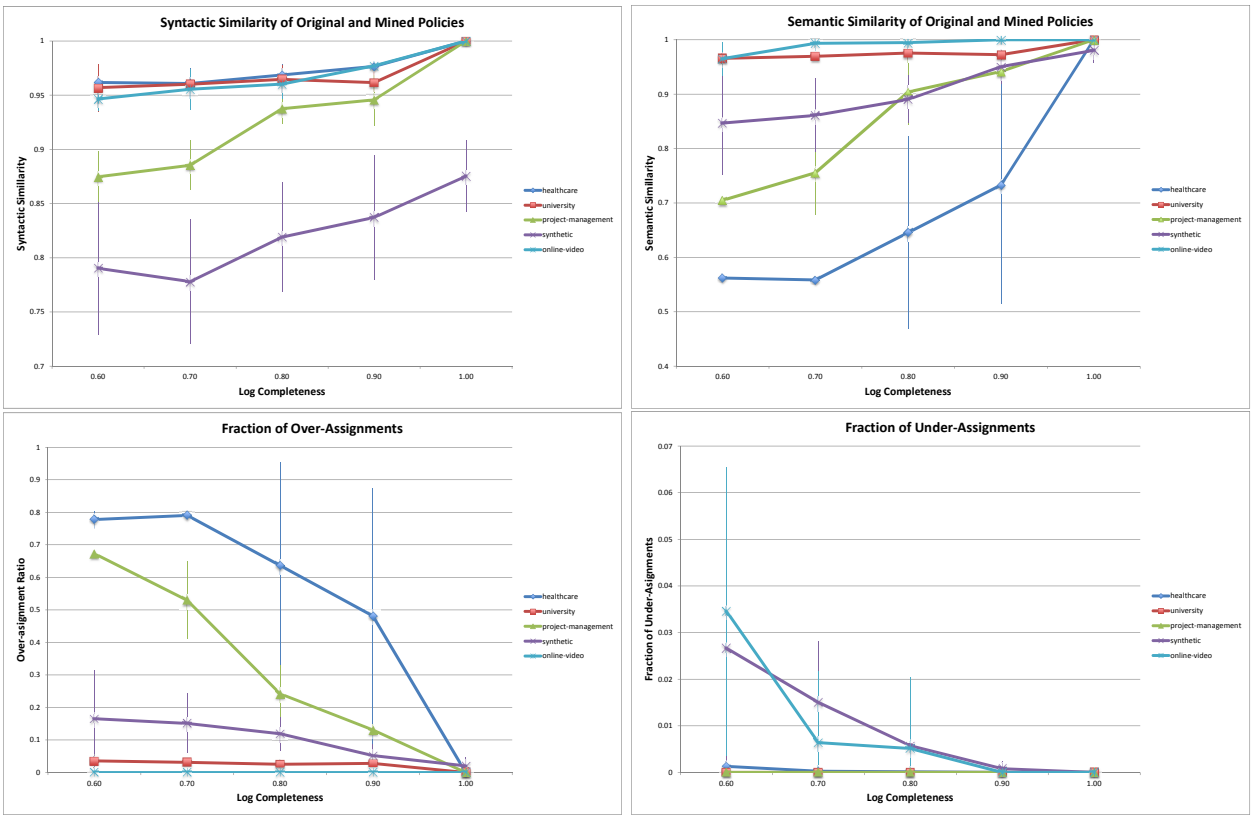

Figure 6: Top: Syntactic similarity and semantic similarity of the original and mined ABAC policies, as a function of log completeness. Bottom: Fractions of over-assignments and under-assignments, as a function of log completeness.

this policy. In fact, if the optimal value of w_o is used for each log completeness, the semantic similarity for healthcare is always above 0.99. Better automated tuning of w_o is a direction for future work.

The fractions of over-assignments and under-assignments are below 0.24 and 0.05, respectively, when log completeness is 80% or higher, for synthetic policies and for case studies other than healthcare. The fractions are higher for healthcare, because w_o is not well chosen, as discussed above.

Comparison of Rule Quality Metrics. The above experiments use the first rule quality metric, Q_{rul}, in Section 4. We also performed experiments using $Q_{\text{rul}}^{\text{freq}}$ and $Q_{\text{rul}}^{\text{ILP}}$ on case studies with manually written attribute data and synthetic policies, and found that Q_{rul} provides the best results overall.

Comparison with Inductive Logic Programming. To translate ABAC policy mining from logs to Progol [Mug97, MF01], we used the translation of ABAC policy mining from ACLs to Progol in [XS13a, Sections 5.5, 16], except negative examples corresponding to absent user-permission tuples are omitted from the generated program, and the statement `set(posonly)?` is included, telling Progol to use its algorithm for learning from positive examples. For the four case studies with manually written attribute data (in contrast, Figure 6 uses synthetic attribute for three of the case studies), for log completeness 100%, semantic similarity of the original and Progol-mined policies ranges from 0.37 for project management and healthcare to 0.93 for online video, while our algorithm exactly reconstructs all four policies. We used Progol 4.4 for these

experiments, because Progol 5.0 segmentation-faults whenever the `posonly` option is set. We also ran the experiments without the `posonly` option using Progol 4.4 and Progol 5.0, and we obtained the same results.

8 Related Work

We are not aware of prior work on ABAC mining from logs. The closest topics of related work are ABAC mining from ACLS, and role mining from logs.

8.1 ABAC Policy Mining from ACLs

Our policy mining algorithm is based on our algorithm for ABAC policy mining from ACLs [XS13a]. The main differences are described in Section 1. The fact that our algorithm can be adapted to work well for this new problem, without changing the overall structure of the algorithm, demonstrates the power and flexibility of our approach.

Ni *et al.* investigated the use of machine learning algorithms for security policy mining [NLC+09]. In the most closely related part of their work, a supervised machine learning algorithm is used to learn classifiers that associate users with roles, given as input the users, the roles, user attribute data, and the user-role assignment. The resulting classifiers are analogous to attribute expressions, but there are several differences between their work and ABAC policy mining. Perhaps the largest difference is that their approach needs to be given the roles and the role-permission or user-role assignment as training data; in contrast, ABAC policy mining algorithms do not require any part of the desired high-level policy to be given as input. Another significant difference is that their work does not consider anything analogous to constraints. Also, they do not consider the size or readability of the generated policy; the classifiers they learn are represented as SVMs.

Association rule mining is another possible basis for ABAC policy mining. However, association rule mining algorithms are not well suited to ABAC policy mining, because they are designed to find rules that are probabilistic in nature [AS94] and are supported by statistically strong evidence. They are not designed to produce a set of rules that are strictly satisfied, that completely cover the input data, and are minimum-sized among such sets of rules. Consequently, unlike our algorithm, they do not give preference to smaller rules or rules with less overlap (to reduce overall policy size).

8.2 Role Mining from Logs

Molloy *et al.* apply a machine learning algorithm that uses a statistical approach, based upon a generative model, to find the policy that is most likely to generate the behavior (usage of permissions) observed in the logs [MPC12]. They given an algorithm, based on Rosen-Zvi et al.'s algorithm for learning Author-Topic Models (ATMs) [RZCG+10], to mine meaningful roles from logs and attribute data, i.e., roles such that the user-role assignment is statistically correlated with user attributes. A discretization algorithm based on simulated annealing is used to produce an RBAC policy from the probabilistic results of the ATM algorithm. The algorithm can be adapted to ABAC policy mining from logs, using a translation from ABAC policy mining to author-topic model construction in which (1) topics correspond to rules, (2) authors correspond to tuples containing a user attribute expression, a resource attribute expression, and a constraint that can appear together in a candidate ABAC rule, (3) documents correspond to a user-resource pair, and (4) words correspond to operations. There is a very large number of such tuples that are well-formed. This is likely to impact the efficiency and effectiveness of the ATM algorithm. If it does not perform well, it might be possible to develop a heuristic to eliminate some authors from consideration before applying the ATM algorithm. In contrast, note that our algorithm never enumerates all well-formed candidate rules; it constructs candidate rules by generalizing specific rules constructed from seeds in the input. Also, note that the number of authors

is much smaller in Molloy *et al.*'s work, which does not consider resource attributes, and hence an author corresponds to only a candidate user attribute expression. Another issue with the ATM algorithm is that it requires specifying the desired number of rules (i.e., topics), which is difficult to predict. Our implementation of the above translation is almost finished. We expect to have experimental results soon.

Zhang *et al.* apply the experience-based access management approach [GLM11] for role evolution [ZGL+11]. In particular, they use machine learning algorithms based on Näive Bayes classifiers to assess the quality of the roles in a given role hierarchy relative to given access logs. The quality is based on the strength of the correlation between the user-role assignment and users' access patterns. The analysis results can be used to improve the role hierarchy. Zhang *et al.* also propose an algorithm based on Support Vector Machines to evolve a role hierarchy [ZCG+13]. Given a role hierarchy and access logs, the algorithm tries to generate a new role hierarchy that optimizes a policy quality metric based on both the distance between the given policy and the generated policy and the homogeneity of the generated roles with respect to the logs. Role homogeneity is a measure of the similarity of the access patterns of the members of a role. These papers do not consider improvement or mining of ABAC policies.

9 Conclusion

This paper presents an algorithm for mining ABAC policies from logs and attribute data. Experiments with case studies and synthetic policies demonstrate the algorithm's effectiveness even when the log reflects only a fraction of the entitlements. Although the original (desired) ABAC policy is not reconstructed perfectly from the log, the mined policy is sufficiently similar to it that the mined policy would be very useful as a starting point for policy administrators tasked with developing that ABAC policy.

Directions for future work include better automated tuning of parameters such as w'_o, characterization of the algorithm's effectiveness as a function of log and policy characteristics such as the ratio of frequencies of the most and least frequent items (*cf.* Section 6.2) and the complexity of the ABAC policy (average number of conjuncts per rule, etc.), and experiments with noisy logs.

Acknowledgments. Christian Hesselbach did the comparison with Progol.

References

[AS94] Rakesh Agrawal and Ramakrishnan Srikant. Fast algorithms for mining association rules. In *Proc. 20th International Conference on Very Large Data Bases (VLDB)*, pages 487–499, 1994.

[BM13] Matthias Beckerle and Leonardo A. Martucci. Formal definitions for usable access control rule sets from goals to metrics. In *Proceedings of the Ninth Symposium on Usable Privacy and Security (SOUPS)*, pages 2:1–2:11. ACM, 2013.

[CDPO08] Alessandro Colantonio, Roberto Di Pietro, and Alberto Ocello. A cost-driven approach to role engineering. In *Proceedings of the 2008 ACM Symposium on Applied Computing (SAC)*, pages 2129–2136. ACM, 2008.

[CDPOV09] Alessandro Colantonio, Roberto Di Pietro, Alberto Ocello, and Nino Vincenzo Verde. A formal framework to elicit roles with business meaning in RBAC systems. In *Proc. 14th ACM Symposium on Access Control Models and Technologies (SACMAT)*, pages 85–94, 2009.

[CDV12] Alessandro Colantonio, Roberto Di Pietro, and Nino Vincenzo Verde. A business-driven decomposition methodology for role mining. *Computers & Security*, 31(7):844–855, October 2012.

[FBB08] Mario Frank, David A. Basin, and Joachim M. Buhmann. A class of probabilistic models for role engineering. In *Proceedings of the 2008 ACM Conference on Computer and Communications Security*, pages 299–310. ACM, 2008.

[FBB10] Mario Frank, Joachim M. Buhmann, and David A. Basin. On the definition of role mining. In *Proceedings of the 15th ACM Symposium on Access Control Models and Technologies*, pages 35–44. ACM, 2010.

[FBB13] Mario Frank, Joachim M. Buhmann, and David A. Basin. Role mining with probabilistic models. *ACM Transactions on Information and System Security*, 15(4):1–28, 2013.

[Fed11] Federal Chief Information Officer Council. Federal identity credential and access management (ficam) roadmap and implementation guidance, version 2.0, December 2011. `http://www.idmanagement.gov/documents/ficam-roadmap-and-implementation-guidance`.

[GLM11] Carl A. Gunter, David M. Liebovitz, and Bradley Malin. Experience-based access management: A life-cycle framework for identity and access management systems. *IEEE Security & Privacy*, 9(5):48–55, 2011.

[GVA08] Qi Guo, Jaideep Vaidya, and Vijayalakshmi Atluri. The role hierarchy mining problem: Discovery of optimal role hierarchies. In *Proc. 2008 Annual Computer Security Applications Conference (ACSAC)*, pages 237–246. IEEE Computer Society, 2008.

[HCBC12] Safaà Hachana, Nora Cuppens-Boulahia, and Frdric Cuppens. Role mining to assist authorization governance: How far have we gone? *International Journal of Secure Software Engineering*, 3(4):45–64, October-December 2012.

[HFK+13] Vincent C. Hu, David Ferraiolo, Rick Kuhn, Arthur R. Friedman, Alan J. Lang, Margaret M. Cogdell, Adam Schnitzer, Kenneth Sandlin, Robert Miller, and Karen Scarfone. Guide to Attribute Based Access Control (ABAC) definition and considerations (final draft). NIST Special Publication 800-162, National Institute of Standards and Technology, September 2013. `http://csrc.nist.gov/publications/drafts/800-162/sp800_162_draft.pdf`.

[KSS03] Martin Kuhlmann, Dalia Shohat, and Gerhard Schimpf. Role mining—revealing business roles for security administration using data mining technology. In *Proceedings of the Eighth ACM Symposium on Access Control Models and Technologies (SACMAT)*, pages 179–186. ACM, 2003.

[LVA08] Haibing Lu, Jaideep Vaidya, and Vijayalakshmi Atluri. Optimal Boolean matrix decomposition: Application to role engineering. In *Proc. 24th International Conference on Data Engineering (ICDE)*, pages 297–306. IEEE, 2008.

[LVAH12] Haibing Lu, Jaideep Vaidya, Vijayalakshmi Atluri, and Yuan Hong. Constraint-aware role mining via extended Boolean matrix decomposition. *IEEE Trans. Dependable Sec. Comput.*, 9(5):655–669, 2012.

[MB00] Stephen Muggleton and Christopher H. Bryant. Theory completion using inverse entailment. In James Cussens and Alan M. Frisch, editors, *Proc. 10th International Conference on Inductive Logic Programming (ILP)*, pages 130–146. Springer, 2000.

[MCL+10] Ian Molloy, Hong Chen, Tiancheng Li, Qihua Wang, Ninghui Li, Elisa Bertino, Seraphin B. Calo, and Jorge Lobo. Mining roles with multiple objectives. *ACM Trans. Inf. Syst. Secur.*, 13(4), 2010.

[MF01] Stephen H. Muggleton and John Firth. CProgol4.4: a tutorial introduction. In S. Dzeroski and N. Lavrac, editors, *Relational Data Mining*, pages 160–188. Springer-Verlag, 2001.

[MLL+09] Ian Molloy, Ninghui Li, Tiancheng Li, Ziqing Mao, Qihua Wang, and Jorge Lobo. Evaluating role mining algorithms. In *Proc. 14th ACM Symposium on Access Control Models and Technologies (SACMAT)*, pages 95–104. ACM, 2009.

[MLL10] Xiaopu Ma, Ruixuan Li, and Zhengding Lu. Role mining based on weights. In *Proceedings of the 15th ACM Symposium on Access Control Models and Technologies*, pages 65–74, 2010.

[MLQ+10] Ian Molloy, Ninghui Li, Yuan (Alan) Qi, Jorge Lobo, and Luke Dickens. Mining roles with noisy data. In *Proceeding of the 15th ACM Symposium on Access Control Models and Technologies (SACMAT)*, pages 45–54. ACM, 2010.

[MPC12] Ian Molloy, Youngja Park, and Suresh Chari. Generative models for access control policies: applications to role mining over logs with attribution. In *Proc. 17th ACM Symposium on Access Control Models and Technologies (SACMAT)*, pages 45–56. ACM, 2012.

[MSAV13] Barsha Mitra, Shamik Sural, Vijayalakshmi Atluri, and Jaideep Vaidya. Toward mining of temporal roles. In *Proc. 27th Annual IFIP WG 11.3 Conference on Data and Applications Security and Privacy (DBSec)*, volume 7964 of *Lecture Notes in Computer Science*, pages 65–80. Springer, 2013.

[Mug95] Stephen H. Muggleton. Inverse entailment and progol. *New Generation Computing*, 13:245–286, 1995.

[Mug97] Stephen Muggleton. Learning from positive data. In *Proceedings of the 6th International Workshop on Inductive Logic Programming (ILP '96)*, volume 1314 of *Lecture Notes in Computer Science*, pages 358–376. Springer, 1997.

[NLC+09] Qun Ni, Jorge Lobo, Seraphin Calo, Pankaj Rohatgi, and Elisa Bertino. Automating role-based provisioning by learning from examples. In *Proc. 14th ACM Symposium on Access Control Models and Technologies (SACMAT)*, pages 75–84. ACM, 2009.

[RZCG+10] Michal Rosen-Zvi, Chaitanya Chemudugunta, Thomas Griffiths, Padhraic Smyth, and Mark Steyvers. Learning author-topic models from text corpora. *ACM Transactions on Information Systems*, 28(1):4:1–4:38, January 2010.

[SS05] Jürgen Schlegelmilch and Ulrike Steffens. Role mining with ORCA. In *Proceedings of the 10th ACM Symposium on Access Control Models and Technologies*, pages 168–176. ACM, 2005.

[TJ10] Hassan Takabi and James B. D. Joshi. StateMiner: an efficient similarity-based approach for optimal mining of role hierarchy. In *Proc. 15th ACM Symposium on Access Control Models and Technologies (SACMAT)*, pages 55–64. ACM, 2010.

[VAG07] Jaideep Vaidya, Vijayalakshmi Atluri, and Qi Guo. The role mining problem: finding a minimal descriptive set of roles. In *Proc. 12th ACM Symposium on Access Control Models and Technologies (SACMAT)*, pages 175–184. ACM, 2007.

[VAGA08] Jaideep Vaidya, Vijayalakshmi Atluri, Qi Guo, and Nabil Adam. Migrating to optimal RBAC with minimal perturbation. In *Proceedings of the 13th ACM Symposium on Access Control Models and Technologies (SACMAT)*, pages 11–20. ACM, 2008.

[VAW06] Jaideep Vaidya, Vijayalakshmi Atluri, and Janice Warner. RoleMiner: Mining roles using subset enumeration. In *Proc. 13th ACM Conference on Computer and Communications Security (CCS)*, pages 144–153. ACM, 2006.

[VAWG10] Jaideep Vaidya, Vijayalakshmi Atluri, Janice Warner, and Qi Guo. Role engineering via prioritized subset enumeration. *IEEE Trans. Dependable Secur. Comput.*, 7(3):300–314, 2010.

[VVAC12] Nino Vincenzo Verde, Jaideep Vaidya, Vijay Atluri, and Alessandro Colantonio. Role engineering: From theory to practice. In *Proceedings of the Second ACM Conference on Data and Application Security and Privacy (CODASPY)*, pages 181–192. ACM, 2012.

[XAC] eXtensible Access Control Markup Language (XACML). `http://www.oasis-open.org/committees/xacml/`.

[XS12] Zhongyuan Xu and Scott D. Stoller. Algorithms for mining meaningful roles. In *Proc. 17th ACM Symposium on Access Control Models and Technologies (SACMAT)*, pages 57–66. ACM, 2012.

[XS13a] Zhongyuan Xu and Scott D. Stoller. Mining attribute-based access control policies. *Computing Research Repository (CoRR)*, abs/1306.2401, June 2013. Revised January 2014. `http://arxiv.org/abs/1306.2401`.

[XS13b] Zhongyuan Xu and Scott D. Stoller. Mining parameterized role-based policies. In *Proc. Third ACM Conference on Data and Application Security and Privacy (CODASPY)*. ACM, 2013.

[ZCG+13] Wen Zhang, You Chen, Carl A. Gunter, David Liebovitz, and Bradley Malin. Evolving role definitions through permission invocation patterns. In *Proceedings of the ACM Symposium on Access Control Models and Technologies (SACMAT)*. ACM, 2013.

[ZGL+11] Wen Zhang, Carl A Gunter, David Liebovitz, Jian Tian, and Bradley Malin. Role prediction using electronic medical record system audits. In *AMIA Annual Symposium Proceedings*, pages 858–867. American Medical Informatics Association, 2011.

[ZRE07] Dana Zhang, Kotagiri Ramamohanarao, and Tim Ebringer. Role engineering using graph optimisation. In *Proceedings of the 12th ACM Symposium on Access Control Models and Technologies*, pages 139–144, 2007.

A Rule Quality Metric Based On Inductive Logic Programming

We review the theory quality metric used in Progol [MB00, MF01], a well-known ILP system, and then describe our design of a rule quality metric based on it. Progol's IPL algorithm works as follows. At the outermost level, Progol uses a loop that repeatedly generalizes an example to a hypothesized rule and then removes examples which are redundant relative to (i.e., covered by) the new rule, until no examples remain to be generalised. When generalizing an example, Progol uses a metric, called a *compression metric*, to guide construction of the hypotheses. When mining ABAC policies from operation logs, user-permissions tuples in $UP(L)$ are positive examples, and no negative examples are available. Thus, this corresponds to the case of learning from only positive data. When learning from only positive data, Progol's compression metric *pcomp* is defined as follows [Mug95].

$$f_m(H) \quad = \quad c \times 2^{-|H|}(1 - g(H))^m \tag{2}$$

$$
\begin{aligned}
pcomp(H,E) &= \log_2 \frac{f_m(H)}{f_m(E)} \\
&= |E| - |H| - m(\log_2(1-g(E)) \\
&\quad - \log_2(1-g(H))) \\
&\approx |E| - |H| + m\log_2(1-g(H))
\end{aligned} \tag{3}
$$

where E is the set of positive examples, H is the entire theory (i.e., ABAC policy, in our context) being generated, including the part not generated yet, $m = |E|$, $|H|$ is the size of H, measured as the number of bits needed to encode H, and c is a constant chosen so that $\sum_{H \in \mathcal{H}} f_m(H) = 1$, where $\mathcal{H}$ is the set of all candidate theories (note that f, like *pcomp* is a function of H and E, since $m = |E|$, but we adopt Muggleton's notation of $f_m(H)$ instead of using the more straightforward notation $f(H,E)$). Let X be the set of all possible well-formed examples (in our context, X is the set of all user-permission tuples). $g(H)$ (the "generality" of H) is the probability that an element of X, randomly selected following a uniform distribution, satisfies H. f_m and *pcomp* can be regarded as policy quality metrics. The term $2^{-|H|}$ in the definition of f_m causes policies with smaller size to have higher quality, and the term $(1-g(H))^m$ causes policies with larger meaning to have higher quality. In [Mug95], f_m is used to guide the search for rules. Using *pcomp* to guide the search would have the same effect, because in the definition of *pcomp* in equation (3), $f_m(E)$ is a constant, and $\log_2 \frac{f_m(H)}{f_m(E)}$ is a monotonic function of $f_m(H)$, so maximizing f_m is equivalent to maximizing *pcomp*.

A difficulty with using f_m to guide generation of a rule to add to a partly generated theory is that the entire theory H is not yet known. To overcome this difficulty, the quality of the entire theory is estimated by extrapolation. Let C_i denote the i'th rule added to the theory, and let $H_i = \{C_1, \ldots, C_i\}$. Let n denote the number of rules in the entire theory being generated (of course, n is not known until the algorithm terminates). When generating C_i, the policy quality $f_m(H_n)$ is estimated as follows.

$$
c \times 2^{-(\frac{m}{p} \times |C_i|)} \times (1 - \frac{m}{p} * (g(H_i) - g(H_{i-1})))^m \tag{4}
$$

where p is the number of examples in E that are implied by C_i and not by H_{i-1}.

Now we describe how to modify our policy mining algorithm to use f_m as a rule quality metric. In the loop in the top-level pseudocode in Figure 1 that builds the set *Rules* of candidate rules, rule quality is computed using equation 4 with $H_{i-1} = Rules$ and $C_i = \rho$. Specifically, in the definition of generalizeRule in 4, $Q_{\text{rul}}(\rho, uncovUP)$ is replaced with f_m computed using equation 4 with $H_{i-1} = Rules$ and $C_i = \rho$. This closely corresponds to the usage in Progol, although it is slightly different, because some of the candidate rules in *Rules* will not be included in the final set of rules *Rules*$'$.

In the loop in the top-level pseudocode in Figure 1 that builds the set *Rules*$'$ of final rules, we take the same approach, except using *Rules*$'$ instead of *Rules*. Specifically, $Q_{\text{rul}}(\rho, uncovUP)$ is replaced with f_m computed using equation 4 with $H_{i-1} = Rules'$ and $C_i = \rho$.

To compute rule quality after building the set *Rules* of candidate rules and before building the set *Rules*$'$ of final rules, the algorithm is modifying a set of rules, not extending a set of rules, so f_m can be evaluated using equation 2, with *Rules* as an estimate of the entire policy H. Specifically, in the calls to elimConjuncts and elimConstraints from simplifyRules, $Q_{\text{rul}}(\rho'', UP_0)$ is replaced with f_m computed using equation 2 with $H = Rules \setminus \rho \cup \{\rho''\}$, and $Q_{\text{rul}}(\rho_{\text{best}}, UP_0)$ is replaced with f_m computed using equation 2 with $H = Rules \setminus \rho \cup \{\rho_{\text{best}}\}$.